

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 February 2003 (20.02.2003) ✓

PCT

(10) International Publication Number
WO 03/014943 A2

(51) International Patent Classification⁷: **G06F 13/00**

(21) International Application Number: **PCT/GB02/03681**

(22) International Filing Date: **9 August 2002 (09.08.2002)**

(25) Filing Language: **English**

(26) Publication Language: **English**

(30) Priority Data:

60/311,505	10 August 2001 (10.08.2001)	US
0119558.5	10 August 2001 (10.08.2001)	GB
60/366,430	21 March 2002 (21.03.2002)	US
0206719.7	21 March 2002 (21.03.2002)	GB
10/143,537	10 May 2002 (10.05.2002)	US

(71) Applicant (for all designated States except US): **SUN MICROSYSTEMS, INC [US/US];** 901 San Antonio Road, Palo Alto, CA 94303 (US).

(72) Inventors; and

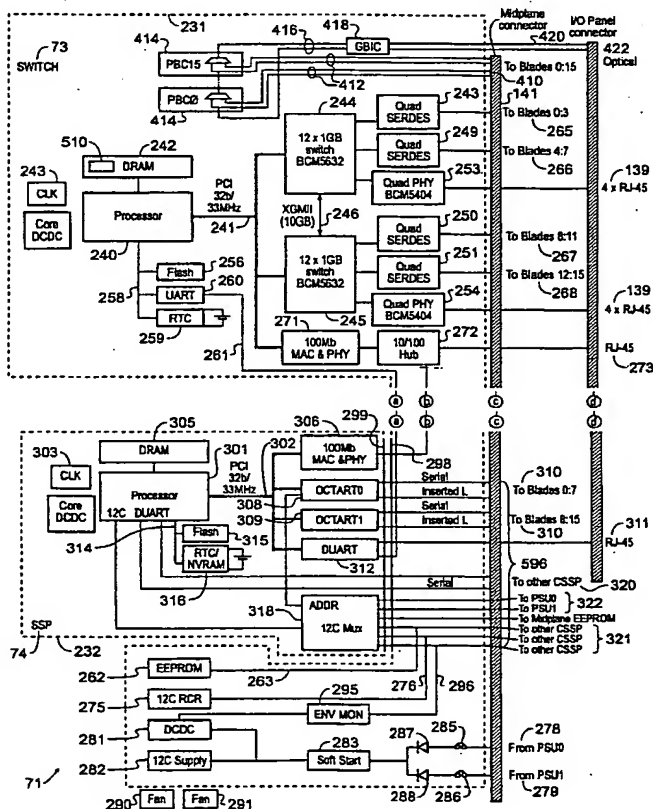
(75) Inventors/Applicants (for US only): **GARNETT, Paul, J. [GB/GB];** Little Orchard, Knightsbridge Road, Camberley, Surrey GU15 3TS (GB). **MAYHEAD, Martin, P. [GB/GB];** Tuscaloosa, Tilford Road, Hindhead, Surrey GU26 6QY (GB).

(74) Agent: **HARRIS, Ian, Richard;** D Young & Co, 21 New Fetter Lane, London EC4A 1DA (GB).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZM, ZW.

[Continued on next page]

(54) Title: **COMPUTER SYSTEM STORAGE**



(57) Abstract: An information processing module, for example a server blade, has information connections and storage connections. A carrier, e.g., a rack mountable shelf, is configured to receive a number of information processing modules. A switch module includes a switch connection for each receiving location that is operable to connect a received information processing module to an external information network. The switch module also includes a port bypass controller for each receiving location that is operable to connect a received information processing module to external storage. Redundant connections are included to enhance reliability.



(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

COMPUTER SYSTEM STORAGE

BACKGROUND OF THE INVENTION

5 The present invention relates to computer systems, in particular to multiprocessor systems, for example multiprocessor server systems.

One application for the present invention relates to high density computer systems, for example, computer server systems for telecommunications applications.

10 In telecommunications applications, it is important to provide high reliability and high capacity of operation. Various approaches have been taken to providing such high-performance, high reliability systems. Typically such systems are designed around providing redundant resources so that if one component of the system develops a fault, the system remains operational using the redundant resources. Fault tolerance

15 can also be achieved, for example, with multiprocessor systems that provide redundancy through dynamic, e.g., software-controlled, task distribution. High density systems are typically rack mountable, with one or more processor systems occupying a shelf in the rack. The trend in recent times is to make the computers with smaller form factors. This means that more computers can be located in a rack. This

20 has the advantage of increasing the processing density within the racks, and also the advantage of reducing the distance between the computer systems.

The present invention relates to the provision of data storage in a higher density computer system.

25

SUMMARY OF THE INVENTION

An aspect of the present invention provides a computer system including at least one information processing module. The information processing module can
5 include a processor, memory, at least one module information connection for communicating information signals and at least one module storage connection for passing data to and from storage. A carrier can include at least one module receiving location configured to receive an information processing module. The module
receiving location can include at least one carrier information connection operable to
10 interconnect with a corresponding module information connection of a received information processing module and at least one carrier storage connection operable to interconnect with a corresponding module storage connection of a received information processing module. The carrier can include at least one switch
connectable to at least one carrier information connection of at least one module
15 receiving location, the switch being operable to connect the information processing module to an external information network. The carrier can also include at least one port bypass controller connectable to at least one carrier storage connection of at least one module receiving location, the port bypass controller being operable to connect the information processing module to external storage.

20

The provision of separate storage connections can enable external storage to be managed efficiently without impacting on the bandwidth of the information connections. The use of a structure of this type enables the connection of direct-attached storage (DAS) and/or connection to a storage area network (SAN) via, for
25 example, Fiber Channel. Redundant sets of connections can be provided to enhance reliability.

The port bypass controllers can be interconnected to form a chain of port bypass controllers, each port bypass controller being connected to a respective carrier storage

connection at each module receiving location. A plurality of information processing modules received in respective receiving locations can thereby be connectable to the external storage.

- 5 The storage connections can be implemented as, for example, Fiber Channel connections. The information connections can be implemented as, for example, Ethernet or Infiniband connections.

In an embodiment of the invention the carrier can be a rack mountable shelf and the
10 information processing unit can be a server blade.

Another aspect of the invention can provide a server blade comprising at least one processor, memory, at least one blade information connection for communicating information signals and at least one blade storage connection for passing data to and
15 from external storage.

The server blade can include a plurality of (e.g., two) information connections for communicating information signals, and a plurality of (e.g., two) storage connections.

- 20 A further aspect of the invention can provide a switch module for a carrier that includes at least one receiving location for receiving an information processing module. A switch connection for each receiving location can be operable to connect a received information processing module to an external information network. A port bypass controller for each receiving location can be operable to connect a received
25 information processing module to external storage.

Another aspect of the invention provides a carrier for a computer system. The carrier can include receiving locations for a plurality of information processing modules, at least one switch module and an interconnection member. The interconnection

member can interconnect information signal channels and storage signal channels between the information processing modules and the switch.

Further aspects and advantages of the invention will become apparent from the
5 following description of particular embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference
5 signs relate to like elements and in which:

Figure 1 is a schematic representation of an architecture of a multiprocessor system for supporting a web site;

10 Figure 2 is a schematic representation of a racking system incorporating an example of a carrier in the form of a rack-mountable shelf according to a first example;

Figure 3 is a front view of an example of a carrier in the form of a shelf of
15 Figure 2;

Figure 4 is a rear view of an example of the shelf of Figure 2;

Figure 5 is a schematic perspective view, partly from the rear, of an example
20 of an information processing cartridge for mounting in the shelf of Figure 2;

Figure 6 is a schematic perspective view of an example of a combined switch and service processor module for mounting in the shelf of Figure 2;

25 Figure 7 is a schematic perspective view of an example of a power supply module for mounting in the shelf of Figure 2;

Figures 8A, 8B and 8C are a schematic plan view and schematic perspective views, respectively, of an example of the chassis and midplane of the shelf of Figure 2;

5 Figures 9A, 9B and 9C are schematic front, top and rear views, respectively, of an example of a midplane of the shelf of Figure 2;

Figure 10 is a schematic cross section view through the shelf of Figure 2;

10 Figure 11 is a functional block diagram of an example of an information processing subsystem for the information processing cartridge of Figure 5;

Figure 12 is a functional block diagram of an example of an information processing subsystem for the combined switch and service processor module of Figure
15 6;

Figure 13 is a functional block diagram of an example of a subsystem for the power supply unit of Figure 7;

20 Figure 14 is a functional block diagram showing the connectivity between the components of the shelf of Figure 2;

Figure 15 is a functional block diagram showing the external connectivity of the shelf of Figure 2;
25

Figure 16 is a schematic representation of a shelf showing the external connections from the shelf of Figure 2;

Figure 17 is a schematic representation of a rack mounted system comprising a plurality of such shelves;

Figure 18 is a functional block diagram of a further example of an information
5 processing subsystem for the information processing cartridge of Figure 5;

Figure 19 is a functional block diagram of an example of combined switch and service processor module subsystem according to the present invention; and

10 Figure 20 is a functional block diagram of an information processing cartridge subsystem according to the present invention.

DESCRIPTION OF PARTICULAR EMBODIMENTS

Embodiments and examples are described hereafter by way of example only in the following with reference to the accompanying drawings.

5

Shown in Figure 1 is an example of an application of a high capacity multiserver system 1 for implementing a network-connected web site such as, for example, an airline reservation system on the World Wide Web.

10 As shown in Figure 1, an external network 3 (e.g., the Internet) for communicating with a user 5 can be connected to gateways 7 which can be connected to an entry edge server group 9 implemented by a web farm as shown in Figure 1. The entry edge server group 9 forms an interface to the external network 3. The entry edge server group 9 can then be connected by switches 11 and a firewall 13 to a web
15 edge server group 15 that can also be implemented as a web farm as shown in Figure 1. The web edge server group 15 can serve to cache web pages that are readily accessible to users 5 accessing the system 1 from the external network 3, for example for checking flight times, etc. The web edge server group can comprise a number of blade server (BS) shelves and a number of network addressable storage (NAS) shelves
20 for storing critical data. The web edge server group 15 can be further connected by a further firewall 17 to a plurality of application servers 19, which can be responsible for, for example, processing flight reservations. The application servers 19 can then be connected via a further firewall 21 to computer systems 23, 25, for example, e-commerce services including financial services for receiving and processing payment
25 for airline reservations.

As will be appreciated, the server system described above with reference to Figure 1 is only an example of a possible application for a multiprocessor server system. Multiprocessor server systems have many different applications and the

present system is not limited to being applicable for use in only one or a limited number of such applications, rather multiprocessor server systems as described herein are operable for use in many different applications. A non-exhaustive list of such alternative applications includes: e-commerce web server systems; 5 telecommunications network server systems; LAN application- and file- server systems and remote vehicle control systems.

With reference to Figure 2, there is shown a schematic perspective representation of a rack system 31 as viewed from the front including left and right 10 front uprights 32 and 33 and left and right rear uprights 34 and 35. The uprights can be formed with apertures for receiving shelf fixings (e.g., screws, bolts, clips, etc., for mounting brackets, slides, rails, etc.).

Also shown in Figure 2 is an example of a blade server shelf 41 mounted in 15 the rack system 31. The shelf 41 forms a carrier configured to carry a plurality of information processing cartridges 43 located side by side along the shelf.

The term "shelf" is used herein in a conventional way to describe a structure that is mountable in rack system 31 and is configured to carry one or more 20 components to form at least a part of a rack-mountable system. In the present example, the shelf 41 is three-dimensional, having a height (H), width (W) and depth (D). In the present example, one dimension (hereinafter described as the height, H) is smaller than the other dimensions (hereinafter described as the depth, D, and the width, W) to facilitate mounting of the shelf in within the rack system 31. It will be 25 appreciated that although the width and depth are typically constrained by the dimensions of the racking system for which the shelf is designed, there is more freedom as regard the height, subject to taking account of any appropriate standards and packaging considerations.

Each of the information processing cartridges contains at least one processor. Each information processing cartridge in the present example is operable as a server. In the described examples, the information processing cartridges are configured as robust enclosed modules.

5

In the example to be described in more detail presently, the information processing cartridges, when aligned in the carrier shelf, look like rectangular slabs, or blades. Accordingly, an information processing cartridge can be described as a blade. The information processing cartridges 43 comprise information processing modules enclosed in an enclosure, or housing, so that the information processing modules have the form of cartridges. Also, as the information processing cartridges are to operate as computer servers in the example described in more detail presently, an information processing cartridge 43 can also be described as a server blade. Accordingly, in the context of this example, the terms module, cartridge and blade are used interchangeably.

The illustrated example of a shelf 41 is configured to carry sixteen information processing cartridges 43, each of which is removably mountable in a respective aperture 45 in the front of shelf, whereby the information processing cartridges can be inserted into and removed from the front of the shelf 41 without removing the shelf 41 from the rack system 31.

In the present example, the shelf 41 comprises a three-dimensional, generally rectangular, enclosure, or housing, 47 that is suitable for mounting in generic racking systems including both 4-post and 2-post systems. It can be mounted on fixed rigid rack mounting ears and/or a simple slide/support system. The present example is designed for standard 19"-wide racking (1"=25.4mm) as defined, for example, in the well-known IEC297 and EIA 310 specification standards with height corresponding to the so-called 3U (3 standard unit) height. For mounting such a 3U unit in such a 19"-

wide racking system, with a depth of, say 25" or 30", the enclosure can be arranged with a height of up to about 130.5mm, a width of up to about 445 mm and a depth, including all hardware and fascias, but excluding cable management, of up to about 635mm, with the depth from the front-most point of a fascia to a rear I/O connector panel of a rear mounted Field Replaceable Unit (FRU) of about 610mm. Of course, other examples designed for other racking systems could have different dimensions.

This example of a shelf 41 has a single enclosure, or housing, 47 that houses a number of modular units or subsystems, the majority of which are replaceable in the field and are therefore known as Field Replaceable Units (FRUs). These modular units include the information processing cartridges 43.

The shelf enclosure 47 can be fabricated from sheet material (e.g., from steel sheet) to form a chassis portion 49 that includes a base 51, two sides 53 and 55, a front 57 and a rear 59. The word "front" as used here is merely used as a label herein to refer to the face, or wall 57 of the enclosure that is located at the main access side of the rack system 31 in use when the shelf is mounted therein. Similarly, the words "rear" and "side" are merely used as labels herein to refer to the faces, or walls 59, 53 and 55 that, in use, are located at those respective positions when the shelf is mounted in the rack system 31.

The apertures 45 can be formed in the front face 57 for receiving the information processing cartridges 43 and, as will be explained later, apertures can also be formed in the rear face 59 for receiving further FRUs. The enclosure can further include a removable top cover 61 that can be secured to the chassis portion 49 by suitable fastening (e.g., screws). The apertures in the front and rear faces 57 and 59 allow at least some of the FRUs to be inserted into and/or removed from the shelf enclosure 47 via the front or the rear thereof, as appropriate, without removing the shelf from the racking. Access to components mounted in the shelf that are not

accessible via one of the apertures in the front 47 and rear 59 faces can be achieved by removing the shelf enclosure 47 from the racking system 31 and then removing the top cover 61 of the shelf enclosure 47.

5 Figure 3 is a front view of an example of a shelf 41 for a first example. A plastic front bezel 63 can be provided that fits on the front face 57 (shown in Figure 2) of the chassis 49 of the shelf enclosure 47. The front bezel 63 can be formed as a unitary removable part that spans the whole width and height of the front of the shelf enclosure 47. The front bezel can include a peripheral portion 64 that can provide
10 areas for corporate and product branding marks, for identification and numbering for the information processing cartridge and for a bar code label (all not shown). One or more apertures 65 can be formed in the peripheral portion 64 of the bezel 63. The apertures 65 in the bezel can be arranged to align with one or more apertures (e.g. a slot (not shown in Figure 3) in the front face of the chassis. In use, air can pass
15 through the apertures 65 to flow into the shelf enclosure 47 to reach FRUs that are mounted in the shelf enclosure 47 through the rear face 59 thereof. Air flowing through the aperture 65 flows into a plenum chamber 66 (not shown in Figure 3) to flow past the processing cartridges 43 to reach rear mounted FRUs. A central area 67 of the front bezel 63 can be open allowing access to the apertures 45 in the front face
20 57 of the shelf enclosure 47 for insertion and removal of the information processing cartridges 43. Where no active module is mounted in a location for an information processing module, a blanking panel, or filler panel, such as the blanking panel 44, can be located in that location. LED indicators 69 can be mounted on a system indicator printed circuit board (not shown) behind a designated area of the bezel to
25 provide an indication of system status via light guides incorporated into the bezel. A further system indicator board (also not shown) carrying LED indicators can be provided inside the shelf enclosure to be visible from the rear thereof.

As mentioned above, in the present example of a shelf, up to sixteen information processing cartridges 43 can be installed in respective apertures 45 in the front face 57 thereof. The number of information processing cartridges 43 actually installed in any installation is dependent upon the system configuration required.

5 Various features relating to the information processing cartridges 43 that are shown in Figure 3 will be described later.

Figure 4 illustrates the rear of the shelf unit of Figures 2 and 3. This shows two different types of FRU 71 and 81 (4 units in total) that have been inserted into

10 respective apertures 72 and 82 in the rear of the shelf enclosure 47. The FRUs shown in Figure 4 include two Combined Switch and Service Processors (CSSPs) 71 and two Power Supply Units (PSUs) 81. Various features shown in Figure 4 will be described later.

15 Before proceeding with a more detailed description of each of the FRUs 43, 71 and 81 introduced so far and of the construction of the shelf 41, there follows a brief description of an information processing cartridge 43, a CSSP 71 and a PSU 81 with reference to Figures 3, 4, 5, 6 and 7.

20 Figure 5 provides a perspective view, partly from the rear, of an information processing cartridge 43. Here it is to be noted that the term "rear" is applied in the context of the position, when installed, of the information processing cartridge, with respect to the shelf 41 (i.e. in this case the "rear" of the information processing cartridge 43 is innermost part of the information processing cartridge when it is

25 inserted in the shelf 41).

With reference to Figures 3 and 5, it will be noted that information processing cartridges are three-dimensional, having a height (h), width (w) and depth (d). If, as in the present example, the information processing cartridges are to be arranged in a

one-dimensional array (a row) across the shelf, then efficient packing for the information processing cartridges is achieved where one dimension (here the width, w) is smaller than the other dimensions (here the depth, d, and the height, h).

5 It will be noted that an enclosure 101 of the present example of an information processing cartridge 43 has six, generally rectangular, faces. For convenience only, the face that is visible from the front of the racking when an information processing cartridge 43 is mounted in the shelf 41 is known as the front face 102. The opposite face is known as the rear face 103. In the present example these two faces, as well as
10 top and bottom faces 104 and 105, have the shape of elongate rectangles. The side faces 106 and 107 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces.

 Although in the present example the information processing cartridges have
15 six generally rectangular faces, it will be appreciated that other examples could have other configurations. For example, rather than having a generally rectangular shape, the side faces of another example of an information processing cartridge could have the general shape of a triangle (whereby the information processing cartridge may then only have five faces), a pentagon (whereby the information processing cartridge may
20 then have seven faces), and so on. Indeed, one or more or all of the edges could be curved. However, it will be appreciated that the present configuration provides advantages for example, in terms of manufacturing, engineering and packing density within a shelf 41.

25 In this example, the information processing cartridge enclosure 101 is fabricated from pressed steel to form two chassis portions. The first portion 109 includes one side face 107, the part of each of the front and rear faces 102 and 103 and the top and bottom faces 104 and 105. The second portion 110 includes the other side face 108 and the remaining part of each of the front and rear faces 102 and 103 and

the top and bottom faces 104 and 105. The two chassis portions 109 and 110 meet at a groove 108 and are secured to one another by fixings (e.g., one or more screws, not shown). The groove 108 runs along the top and bottom faces 104 and 105 of the enclosure 101 and are provided for interfacing with guide rails of the shelf chassis 49 (not shown in Figure 5, but see Figure 8C). A cover portion that is secured to the chassis portion forms the other side face 106. It will be appreciated however, that in another example, the chassis portions could be joined at a position other than the groove 108, with the groove 108 being formed entirely in one of the chassis portions. Alternatively, the enclosure 101 may be constructed from a number of sheets of steel, with each sheet forming one of the faces. The provision of the enclosure means 101 that the information processing cartridge 43 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 101, the information processing cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the information processing cartridge includes its own electromagnetic shielding. To this end, the enclosure 101 is also provided with electromagnetic interference (EMI) fingers 114 to ensure good contact with the shelf chassis and the adjacent components.

As shown in Figure 3, the information processing cartridge 43 incorporates an injector/ejector handle 111 on the front face 102, with an associated latch mechanism 113 shown in Figure 5 for facilitating insertion and latching of the information processing cartridge 43 within an aperture in the shelf 41. As shown in Figure 3, the handle 111 of the injector/ejector lever extends substantially the whole height of the front face of the information processing cartridge 43, thereby increasing the mechanical advantage and facilitating injection and ejection of the information processing cartridge 43. As is further shown in Figure 3, the front face 102 of the information processing cartridge 43 has perforations 115, in the present example slits, to allow for airflow into the information processing cartridge 43. It will be noted in

Figure 3 that the handle 111 is narrower in its middle than at its ends. This reduces any masking effect of the handle 111 on the airflow to the perforations 115 in the front face of the information processing cartridge 43 and facilitates grasping of handle 111.

The handle 111 can be bowed out from the front face of the information processing
5 cartridge to further facilitate grasping thereof and to reduce any masking effect with regard to the perforations. In other examples, the handle 111 could have an open frame-like structure to further facilitate airflow. As shown in Figure 5, the rear face 103 of the information processing cartridge 43 also has perforations 117 to allow for air to be exhausted from the rear of the information processing cartridge 43. A fan
10 can be located within the enclosure 101 of an information processing cartridge 43. In the present example the fan is an impingement fan to direct cooling air onto a processor of the information processing cartridge, which fan also encourages air movement within the enclosure 101. LED indicators 119 (see Figure 3) can be provided on the front face 102 of an information processing cartridge to indicate
15 whether power is on, whether service intervention is required and whether the information processing cartridge 43 can be removed. A connector 120, for example a 40 way single connector attach (SCA-2) connector, can be provided at the rear of the information processing cartridge 43 for electrical connection of the information processing cartridge 43 within the shelf 41. The connector 120 is advantageously able
20 to withstand repeated removals and insertions of the enclosure 101 from and into a shelf 41. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the information processing cartridge into the receiving location.

25 Figure 6 provides a perspective view, partly from the front, of a Combined Shelf and Service Processor (CSSP) cartridge 71. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the CSSP cartridge 71, with respect to the shelf 41 (i.e. in this case the "front" of the CSSP cartridge 71 is innermost part of the CSSP cartridge 71 when it is inserted in the shelf 41).

With reference to Figures 4 and 6, it will be noted that a CSSP cartridge 71 is three-dimensional, having a height (h), width (w) and depth (d). An enclosure 121 of present example of a CSSP cartridge 71 has six, generally rectangular, faces. For convenience only, the face that is visible from the rear of the racking when a CSSP cartridge 71 is mounted in the shelf 41 is known as the rear face 122. The opposite face is known as the front face 123. In the present example these two faces, as well as side faces 126 and 127 have the shape of elongate rectangles. The top and bottom faces 124 and 125 are also rectangular, but not elongate in the manner of the front, rear, top and bottom faces. Although in the present example the CSSP cartridges have six generally rectangular faces, as for the information processing cartridges 43 it will be appreciated that other examples could have other configurations.

In this example, the CSSP cartridge enclosure 121 is fabricated from steel sheet to form a chassis portion that includes the bottom face 125, the front and rear faces 122 and 123 and the side faces 126 and 127. A cover portion that is secured to the chassis portion forms the other top face 124. The cover portion is secured to the chassis portion by suitable fixings, for example one or more screws 128. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 121 means that the CSSP cartridge 71 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 121, the switch cartridge is a robust unit that protects its inner workings from the outside environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the CSSP cartridge includes its own electromagnetic shielding. To this end the CSSP enclosure 121 is provided with EMI fingers 129 to ensure good contact with the shelf chassis and the adjacent components.

As shown in Figure 4, the CSSP cartridge 71 incorporates two D-shaped handles 132 to facilitate insertion and removal of the CSSP cartridge 71 with respect to an aperture 72 in the rear face of the shelf enclosure. A latch member 131 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 131 is configured to engage one of the handles 132 and to secure the CSSP 71 in place. In other examples, the CSSP 71 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 6, the front face 123 of the CSSP cartridge 71 has perforations 133 to allow for airflow into the CSSP cartridge 71. As shown in Figure 4, the rear face 122 of the CSSP cartridge 71 has perforations 135 to allow for air to be exhausted from the rear of the CSSP cartridge 71.

At least one fan can be located, for example behind the perforated portion 135 of the rear face, in a CSSP cartridge 71 to channel cooling air through the CSSP cartridge 71 from the front to the rear. In this particular example shown, two fans are provided, one behind each set of perforations 135. LED indicators 137, as shown in Figure 4, can be provided on the rear face 122 of the CSSP enclosure 121 to indicate whether power is on, whether service intervention is required and whether the switch can be removed. Additional link status indicators can be provided integral to 2x4 stacked RJ-45 connectors 139, also shown in Figure 4. As shown in Figure 4, electrical connections 141 can be provided at the front face of the CSSP (i.e. on the face that in use is inside the shelf enclosure 47). Suitable connections for use in the present example include a connector for power connections, a connector for serial management data connections and a connector for information connections. In the present example, information connections are implemented using an Ethernet information communication protocol, e.g. at 1 Gigabit (Gb). However other protocols could equally be used, for example the Infiniband information communication protocol. The connector arrangement can include a guide pin arrangement to prevent module misalignment during insertion of the CSSP cartridge module into the

receiving location. For this purpose, guide pin holes 142 can be provided on the front face 123 into which guide pins may pass to aid module alignment.

In the present example, up to two CSSPs 71 can be mounted at any one time at
5 the rear of the shelf unit in corresponding apertures 72 in the rear face of the shelf enclosure 47. The number of CSSPs 71 provided in any particular implementation depends upon system configuration and the need, or otherwise, for redundancy.

Figure 7 provides a perspective view, partly from the front, of a power supply
10 unit (PSU) cartridge 81. Here it is to be noted that the term "front" is applied in the context of the position, when installed, of the PSU cartridge 81, with respect to the shelf 41 (i.e. in this case the "front" of the PSU cartridge 81 is innermost part of the PSU cartridge 81 when it is inserted in the shelf 41).

15 With reference to Figures 4 and 7, it will be noted that PSU cartridge 81 is three-dimensional, having a height (h), width (w) and depth (d). In this particular example, the order to provide for a dense packing of the FRUs in the rear of the shelf 41, the PSU cartridge 81 has two dimensions (hereinafter described as the width, w, and the depth, d) that are generally similar. An enclosure 145 of present example of a
20 PSU cartridge 81 is of generally oblong shape, but has the "top" "front" edge cut away to form an additional "top" "front" sloping face. The enclosure 145 therefore has five, generally rectangular, faces and two faces of generally rectangular shape with one corner cut away. For convenience only, the face that is visible from the rear of the racking when the PSU cartridge 81 is mounted in the shelf 41 is known as the rear
25 face 146. The opposite face is known as the front face 147. In the present example these two faces and the two side faces 150, 151 are of elongate generally rectangular shape with one corner cut away, given that the width and depth of the PSU cartridge are similar, whereas the top and bottom faces 148, 149, although still rectangular, are not, in this example, notably elongate. A top front face 148a is present at the top front

of the enclosure. Thus the front of the enclosure is sloped at the top edge. As for the information processing cartridges 43, however, it will be appreciated that other examples could have other configurations.

5 In this example, the PSU cartridge enclosure 145 is fabricated from steel sheet to form a housing portion that includes the bottom face 149, the side faces 150 and 151 and the front and rear faces 146 and 147. Cover portions that are secured to the housing portion form the top face 148 and top front face 148a. The cover portions are secured to the chassis portion by suitable fixings, for example one or more screws
10 152. It will be appreciated however, that in another example, other faces, or portions, of the enclosure could form the chassis and the cover portions. The provision of the enclosure 145 means that the PSU cartridge 81 can safely be handled by an operator who is not a skilled technician. Also, through the use of the enclosure 145, the PSU cartridge 81 is a robust unit that protects its inner workings from the outside
15 environment and vice versa. The use of a conductive enclosure, e.g., a metal enclosure, means that the PSU cartridge includes its own electromagnetic shielding. To this end the PSU enclosure 145 is provided with EMI fingers 153 to ensure good contact with the shelf chassis and the adjacent components.

20 As shown in Figure 4, the PSU cartridge 81 incorporates two D-shaped handles 156 to facilitate insertion and removal of the PSU cartridge 81 with respect to an aperture 82 in the rear face of the shelf enclosure. A latch member 155 can be pivotably mounted on a plate that can be secured (e.g., using screws) to the rear face of the shelf enclosure. The latch member 155 is configured to engage one of the
25 handles 156 and to secure the PSU 81 in place. In other examples, the PSU 81 could be provided with an injector/ejector handle in a manner similar to the information processing cartridge. As shown in Figure 7, the front face 147 of the PSU cartridge 81 has perforations 157 to allow for airflow into the PSU cartridge 81. As shown in

Figure 4, the rear face 146 of the PSU cartridge 81 also has perforations 159 to allow for air to be exhausted from the rear of the PSU cartridge 81.

A pair of fans can be located behind the perforated portions 159 of the rear face of a PSU cartridge 81 to channel cooling air through the PSU cartridge from the front to the rear. LED indicators 161 can be provided on the rear face 146 of the PSU enclosure 81 to indicate whether input power is good, whether output power is good, whether service intervention is required and whether the PSU can be removed. Electrical connectors 163 can be provided at the front face of the PSU (i.e. on the face that in use is inside the shelf enclosure 47) for connection to the shelf. The PSU 81 of the present example may suitably employ an SSI-MPS (Server Systems Interface - Midrange Power Supply) compliant right angle connector at the front face 147 of the PSU 81 to connect to the shelf 41. The power inlet 83 for each PSU 81 can incorporate a cable/connector retention mechanism (not shown) on the rear face 146 of the PSU to prevent accidental or malicious removal of the power input cord from the PSU 81.

In the present example, the shelf unit enclosure provides slots at the rear of the shelf unit for two hot-swappable, AC input PSUs 81. To provide redundant power supply, both PSU slots are populated. More generally, $N+M$ power supplies can be provided, where N is the minimum number of power supply units needed to support the components mounted in the shelf, and M is a selectable number of 0, 1 or more power supplies to provide for redundancy. In the present example, $N = M = 1$.

If the full complement of information processing cartridges or switches are not fitted to the enclosure, then blanking panels/modules (e.g., the blanking panels 44 shown in Figure 3) are fitted to all unpopulated positions to maintain the correct airflow and thermal characteristics of the shelf, a safe internal operating temperature.

for the shelf life expectancy of all shelf components, electromagnetic compliance (EMC) containment and electrostatic discharge (ESD) containment.

It will be noted that each of the FRUs, such as the information processing
5 cartridges 43, is advantageously contained in its own robust enclosure to facilitate
EMC containment, ESD containment, handling, storage and transportation. Each
FRU can be configured as a 'sealed' unit in the sense that it can be configured not to
have field or customer serviceable parts internally. The FRUs can be configured
readily to plug into the shelf enclosure and to be hot swappable. The FRUs can be
10 keyed to prevent incorrect positioning and insertion into the shelf enclosure and are
arranged positively to be retained in the shelf by a latching/locking mechanism.

The examples of FRUs described above are not provided with removable
media. In the present example, internal data storage is provided by 2.5" IDE 9.5mm
15 or 12.7mm profile hard disk drive (HDD) devices mounted internally in each
information processing cartridge 43 and in the CSSP cartridge 71. The drives are not
considered as FRUs and are not hot-swappable disk drives in the present example,
although they could be in other examples. In other examples, the information
processing cartridges can be configured without internal hard disk drives.
20

The internal configuration of the shelf 41 and a midplane 171 contained
therein is described in the following with reference to Figures 8A, 8B and 8C, and
Figures 9A, 9B and 9C.

25 Figure 8A is a schematic plan view showing the internal configuration of an
example of a shelf 41 with the cover 61 removed. Figure 8B is a schematic
perspective view from above the rear of the chassis portion 47 of the shelf enclosure
with the field replaceable units removed. Figure 8C is a schematic perspective view
from below the front of the chassis portion 47 of the shelf enclosure with the field

replaceable units and the base 51 removed. Figures 9A, 9B and 9C are, respectively, front, top and rear views of the midplane 171. In this example, the midplane is, in use, mounted vertically within the shelf 41 extending across the width W of the shelf 41 at a position approximately half way between the front and the rear of the shelf 41.

5

The vertically mounted midplane 171 extends, in this example, across the shelf 41 and allows for the electrical interconnection of the FRUs. The various apertures in the front and rear faces 57 and 59 of the shelf 41, in combination with the midplane 171, can be provided with guides (e.g., rails 181) and keying e.g., offset connector positioning for the insertion of the FRUs into the enclosure and midplane 171. The midplane 171 can be a double-sided, or multi-layer printed circuit board (PCB) assembly that can be mounted vertically in a rigid manner within the enclosure. It can carry connectors 175 on a front surface 172 for making electrical connection with corresponding connectors 120 on the information processing cartridges 43. It can also carry connectors 177 and 179 on rear surface 173 for making electrical connection with corresponding connectors 141 and 163 on the CSSPs 71 and the PSUs 81, respectively. Conductive tracks (not shown) on and through the midplane 171 can be provided to interconnect the various connectors. In addition, the midplane can provide connectors for receiving corresponding connectors connected to first and second indicator boards 183 and 184 that each carry a respective set of LED indicators 69. In the present example, the midplane 171 is not configured as a FRU and is not hot swappable. It is perforated to facilitate airflow through the shelf 41. The midplane 171 can include openings 185, which cooperate with openings in the enclosures of the FRUs 43, 71 and 81, to provide a path for cooling air to pass from the front to the rear of the shelf 41, the cooling air being forced through by fans in at least selected ones of the FRUs, for example in the CSSPs 71 and the PSUs 81, possibly also in the information processing cartridges 43.

A plenum chamber floor member 94 can extend horizontally from the front of the midplane 171 to the front face 57 of the shelf enclosure, or chassis 47. The member 94 provides a floor for a plenum chamber 66, which is supplied with air via the apertures 65 in the front bezel and, in the illustrated example, the slot shaped aperture 68 in the front face 57 of the shelf enclosure 47. The top and sides of the plenum chamber are provided by the top cover 61 and side faces 53 and 54 of the shelf enclosure 47.

A plurality of cartridge guide rails 98 can be provided at the underside of the plenum chamber floor member 94. These rails can be made from a material having a low coefficient of friction, such as polytetrafluoroethene (PTFE) or polythene. Each rail is advantageously positioned so as to interface with the groove 108 of each processing cartridge 43 to aid correct alignment and to facilitate insertion of the processing cartridge during insertion of the cartridge into the shelf 41. A further row of guide rails (not shown) can be provided in matching positions at the top surface of the base 51 of the shelf 41 such that alignment guide rails are present above and below the processing cartridges.

A CSSP/PSU divider 96 can be provided to the rear of the midplane 171 and can extend horizontally to the rear face 59 of the shelf enclosure 47. The CSSPs 71, when inserted, are supported by the divider 96. To aid the correct insertion of the CSSPs 71, CSSP guide pins 178 are provided on the midplane 171 at positions adjacent connectors 177 on the midplane 171 for connection to the CSSPs 71.

Respective positions 88 and 89 can be formed in the front face 57 and the rear face 59 at which first and second indicator boards 183 and 184 supporting the indicator LEDs 69 can be located. These positions 88, 89 therefore include an aperture through the respective face of the shelf enclosure 47 such that indicator LEDs

69 mounted onto a circuit board attached to the inside of the shelf enclosure 47 may be viewed from outside the shelf enclosure.

There now follows are more detailed description of the midplane 171.

5

As mentioned above, the midplane 171 connects all the elements of a shelf together, including, in the present example, up to sixteen information processing cartridges 43, up to two CSSPs 71, two PSUs 81 and the two indicator boards 183 and 184. In the present example, due to its location within the shelf enclosure, the midplane 171 is not configured to be swappable. Accordingly, to maximize the system reliability, the midplane is configured to provide as a high level of reliability as possible. To this end, the midplane is advantageously configured without active devices and to include the minimum number of decoupling capacitors consistent with good design practice (ideally zero).

15

The midplane supports a number of paths for various power and signal lines to interconnect the FRUs.

In the present example, each information processing cartridge 43 has a high speed information signal connection (e.g., a Gigabit (Gb) Ethernet SERIALiser/DESerialiser (SERDES) connection) to each of the CSSPs 71, each connection consisting of two pairs of differential signals. In a conventional manner therefore, the tracking of the paths for these signals is arranged to keep the pairs well balanced and on a single signal layer (i.e. without vias) to support such differential signals at high frequency.

In addition, in the present example, each information processing cartridge 43 has a serial console connection to the CSSP 71. Each connection consists of two TTL

(Transistor-Transistor Logic) level signals that make a transmit and return (TX and RX) pair.

Also, each PSU 81 has a management signal connection (e.g., a serial I2C (Inter-IC Bus) connection) to the CSSP 71 to control power and monitor environmental parameters. The I2C bus comprises of two signals SCL and SDA (serial clock line and serial data line). In addition, an I2C address programming pin is provided for the PSUs 81.

Each information processing cartridge 43 and PSU 81 can signal to the CSSP 71 that it is inserted by pulling to ground (GND) a respective Inserted_L signal (i.e., an active low signal). These signals are fed to the CSSP 71 via the midplane 171.

Each PSU 81 has five 12 Volt output rails. The routing from each PSU 81 is arranged so that a fault in any single FRU cannot completely interrupt the power to any other.

As mentioned above, the midplane 171 is provided with appropriate connector arrangements for receiving the connectors on the FRUs.

20

In the present example, the information processing cartridge 43 connects to the midplane 171 through a 40 pin Single Connector Attachment (SCA-2) connector as defined by the Small Computer Systems Interface (SCSI) standard. Accordingly, the midplane carries corresponding connectors 175.

25

In the present example, each CSSP 71 connects to the midplane 171 through a right-angle 40 pair connector. The corresponding connectors 177 on the midplane are straight male parts with a power connector. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during

insertion. A power module is used to provide connection between the power rails. The center pin is longer on the switch modules to provide a leading ground. The CSSP 71 also connects to the midplane 171 through a right-angled 125 way 5 row 2mm connector. The connector 177 on the midplane 171 includes a straight male
5 part. A guide pin arrangement is provided in addition to the connectors to prevent misaligned modules causing bent pins during insertion.

In the present example, as mentioned above, each PSU 81 connects to the midplane 171 through an SSI-MPS specification connector. The contacts are
10 configured 5P/24S/6P with sequenced signal (S) and power (P) pins. Where the connector on the PSU is a 1450230-1 R/A male header, solder tails connector, the mating connector 179 on the midplane can be a 1450540-2 vertical receptacle, press-fit connector.

15 In the present implementation, indicator boards 183 and 184 (see Figure 8A) are provided at the front and rear of the system and are configured as FRUs. In this example they hold three system-level indicator LEDs 69 and include a FRU identity (FRU-ID) programmable read-only memory (PROM) each. Three LEDs 69 are present on the indicator board. There can, for example, be a white locator LED that
20 can be switched by the user for locating the system; a green power-on LED to indicate when the system is powered; and a red service-required LED to indicate a fault or other condition requiring servicing. These LEDs can be driven by the CSSP 71.

In the present example, identification information (FRU ID) for the midplane
25 171 is held on an I2C electrically erasable programmable read only memory (EEPROM) in the front indicator board 183. In addition to the I2C signals necessary to access the FRU ID EEPROM, the CSSPs 71 provides a current limited supply to the indicator boards 183 and 184 via the midplane. The indicator boards 183 and 184 are also provided with an I2C address programming pin. Depending on the

implementation, FRU ID information can be stored instead, or in addition, on the rear indicator board 184.

As the FRU-ID for the midplane 171 is held on one or both of the indicator
5 boards 183 and 184, the midplane can be a totally passive unit. The FRU-ID PROMs
communicate with the CSSPs 71 via an I2C bus. Each device on the bus has a
separate I2C address. The lower three I2C address bits of the EEPROMs used are
available as pins on the device, to allow programming with resistors. The least
significant bit of this address (A0) is passed to the midplane via the corresponding
10 connector. This allows the midplane 171 to program the address of the FRU-ID
differently for the front and rear indicator boards 183 and 184, by pulling the address
low for the front board and high for the rear indicator board 183. This ensures that
both EEPROMS are available on the bus, at different addresses. The FRU-ID for the
midplane and for the other FRUs to be installed in the shelf can be stored on either
15 front or rear EEPROM, but the CSSPs 71 are configured to use the EEPROM on the
front indicator board 183 in the present embodiment. The EEPROM is at least 8kByte
in size.

As mentioned above, the midplane 171 includes openings 185 to provide a
20 ventilation path for cooling air passing through the shelf 41. The cooling air is forced
through the shelf 41 by means of fans provided in each of the information processing
cartridges 43, the CSSP modules 71 and the power supply modules 81. The openings
185 shown in Figures 8B, 9A, 9B and 9C form schematic representations of openings
in the midplane 171. In practice, the openings could have any form (i.e., a series of
25 large openings, or a number of small perforations), arranged on the midplane to align
with corresponding openings or ventilation apertures in the various field replaceable
units 43, 71 and 81. In this way, the path of the airflow from the front of the shelf to
the back of the shelf can be configured to be as efficient as possible, depending on the
detail configuration of the fan units and the ventilation openings or apertures in the

- information processing, switch, service processor and power supply unit modules 43, 71 and 81. Providing the fan units in the field replaceable units 43, 71 and 81, contributes to the aim of maintaining the chassis 49 and the midplane 171 of the shelf 41 free of active components, thereby minimising cost, and facilitating maintenance.
- 5 Also, by providing the fan units in each of the field replaceable units, merely inserting and removing field replaceable units automatically adapts the flow of cooling air to the number and type of field replaceable units inserted in the shelf 41.

- As described above, in the present example each of the FRUs is designed to be
- 10 a non-user serviceable unit. Thus each FRU presents the user with a "sealed" unit which may be inserted into and removed from the shelf 41 as desired or required. If a FRU ceases to be operable, then the user has a choice only of returning the FRU to a supplier or service company for repair or of discarding the non-operable unit. As the FRUs are non-user serviceable, there is no requirement for a skilled technician to be
- 15 employed in inserting or removing the FRUs into or from a shelf 41. Thus each FRU is designed such that a non-skilled person should have difficulty in causing damage to the FRU during handling. Moreover, the configuration and construction of the FRUs (e.g., the injector/ejector levers, the grooves in the enclosures of the information processing units, etc), of the shelf enclosure and the midplane (e.g., the guide rails to
- 20 guide insertion of the FRUs, the locating pins, etc) contribute to facilitating easy insertion and removal of the FRUs.

Shown in Figure 10 is an example the flow of cooling air through the shelf 41 and FRUs 43, 71, 81 mounted therein.

25

In this example, the cooling air passing through the shelf 41 is drawn generally in a front to rear direction through the shelf 41 by cooling fans mounted within the CSSPs 71 and the PSUs 81. Two separate flow paths for cooling air are provided in this example. The first, indicated as flow path $\alpha\alpha$ by dotted lines 77 provides cooling

air to the CSSPs 71. The second path, indicated as flow path $\gamma\gamma$ by dotted lines 78 provides cooling air to the information processing cartridges 43 and PSUs 81.

The flow of cooling air along path $\alpha\alpha$ enters the shelf 41 through the elongate
5 aperture 65 in the front face 57 of the shelf enclosure 47. This air then flows through the plenum chamber 66, and passes over the top edge of the midplane 171 to reach the perforations 133 of the front face of the CSSPs 71. The cooling air then passes through the CSSPs 71, providing cooling to the components thereof before passing out of the CSSPs 71 through the perforations 135 in the rear face of the CSSPs 71 thus
10 being exhausted from the shelf 41. This flow of cooling air along flow path $\alpha\alpha$ is driven by fans 79 mounted within the CSSPs 71. In the present example, a pair of fans 79 is provided within each CSSP 71 and are mounted against the rear face thereof.

15 Air flowing along path $\alpha\alpha$ is impeded from flowing around the processing cartridges 43 by plenum chamber floor member 94 and is impeded from flowing to the PSUs 81 by CSSP/PSU divider 96. This flow path $\alpha\alpha$ therefore ensures that air flowing to the CSSPs 71 is not warmed by passage through the processing cartridges 43 and therefore provides maximum efficiency cooling to the CSSPs 71.

20

The flow of cooling air along path $\gamma\gamma$ enters the shelf 41 through the perforations 115 in the front face of the information processing cartridges 43. The air thus enters the information processing cartridges 43 and provides cooling to the components thereof. Cooling fans (not shown) within the information processing
25 cartridges 43 direct the cooling air to the processor (CPU) of the information processing cartridge and directs the flow of air in the cartridge thereby increasing cooling efficiency. The air then exits the information processing cartridges 43 through the perforations 117 in the rear face thereof. The air then passes through the apertures 185 through the midplane 171 to reach the PSUs 81. This cooling air then

passes through the perforations 157 on the front and upper front faces of the PSUs 81 to enter the PSUs and provide cooling to components thereof. It will be appreciated from Figure 10 that the sloping rear of the upper face of the PSUs 81 increases the area over which air can be drawn into the PSUs, thereby reducing the back pressure on the air flowing through the shelf unit and aiding the cooling efficiency. The flow of cooling air along path $\gamma\gamma$ is driven by fans 85 mounted within the PSUs 81. In the present example, a pair of fans 85 is provided within each PSU 81 and are mounted against the rear face thereof.

10 Air reaching the PSUs 81 via path $\gamma\gamma$ will already have passed through the processing cartridges 43. Such air will therefore be already warmed above the ambient temperature outside of the shelf 41 by its passage through the processing cartridges 43. However, as the cooling requirement of the PSUs 81 is typically less than that for the CSSPs 71, this does not cause any difficulty in the operation of the PSUs 81, which are adequately cooled by this flow of pre-warmed air. The pre-warmed air passing through the apertures 185 through the midplane 171 is impeded from flowing into path $\alpha\alpha$ and entering the CSSPs 71 by the SCCP/PSU divider 96.

As will be appreciated by the skilled addressee, the arrangement shown in Figure 10 is illustrative only and other arrangements whereby, for example, a mixture of cool air from the plenum chamber 65 and warmed air from the processing cartridges 43 is provided to each rear-mounted FRU can readily be constructed.

With reference to Figure 11, there now follows a description of functional elements of an information processing cartridge 43 as contained within the information processing cartridge enclosure 101.

The information processing cartridge 43 includes a microprocessor 192 (a non-limiting example of a microprocessor that can be utilised in the present example is an

UltraSPARC™ processor). The microprocessor is mounted on an information processing cartridge motherboard 191.

5 A configurable clock generator 193, configured as a programmable clock synthesizer employing a crystal, can be used to produce CPU clock signals, CLKA and CLKB. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) 194 and a configurable core voltage regulator module (VRM) 195 are provided.

10 In the present example, memory means for use by the processor 192 when executing instructions can be provided in the form of buffered dynamic random access memory (DRAM), for example configured as dual in line memory modules (DIMMs) 196 with a 72-bit data path with error correction codes (ECC), seated in two sockets on a riser card from the information processing cartridge motherboard 191. The
15 memory capacity can be chosen to suit the processor addressable memory space. For example, in the present example, up to 4 Gigabytes (4GB) of addressable memory can be provided. Serial Presence Detect (SPD) auto-configuration is provided via a Service Management Bus (SMBus) over an I2C bus 197.

20 In the present example, a PCI bus architecture can be employed with a so-called SouthBridge bus bridge 199 with SuperIO and two Gb Ethernet Media Access Control (MAC) devices. As described above, however, other bus protocols (e.g., Infiniband) can be used. A 32bit PCI bus 198 can be provided from the microprocessor 192. The SouthBridge 199 is a standard form of bus bridge, in the
25 present example packaged in a 352 pin PBGA (Plastic Ball Grid Array) package, that provides the following functions: an SM Bus interface over the I2C bus 197 for access to the SPD (Serial Presence Detect) feature of the DIMMs that allows initialization of the memory controller; an Xbus interface for access via an Xbus 200 (which is a packet switched multiprocessor bus) to a PROM 201, a real time clock

(RTC) 202 and an information processing cartridge service controller (hereinafter termed a Blade Service Controller (BSC)) 203; an IDE (Integrated Drive Electronics) interface that provides an ATA-100 (AT Attachment) IDE connection 204 to an IDE disk drive 205; and a serial console interface on a service bus 206 to the BSC 203 that
5 is used for operating system functions including a console function with this embodiment.

For IO to the midplane 171, two AC-coupled Ethernet interfaces 207 and 208 are provided in the present example, which are packaged in a 316 pin PBGA. These
10 Ethernet interfaces can provide a PCI attached Ethernet MAC capable of operation up to Gigabit Ethernet performance. The physical layer can be implemented using SERIALiser/DESerialisers (SERDESs) 209 and 210. An example of a SERDES device is the TLK2201 transceiver manufactured by Texas Instruments, Inc. The SERDES devices use differential PECL TX+/- and RX+/- (Positive Emitter Coupled Logic
15 Transmit and Receive) pairs to communicate to the switch portions of the CSSPs 71 over the midplane 171. The RX+/- pairs can be AC coupled at the information processing cartridge 43, the TX+/- pairs can be AC coupled at each CSSP 71. This facilitates hot-swap of the information processing cartridges 43 and the CSSPs 71.

20 Asynchronous serial connections 211 and 212 for communication between the BSC 203 and the Service Processor parts of the CSSPs 71 can be provided.

Internal data storage can be provided in the present example by a hard disk 205 with a capacity of 30GB or more rated for 24/7 continuous operation. The hard disk
25 205 is accessed using the primary IDE interface of the SouthBridge 199. The hard disk 205 can hold an operating system, for example a Solaris operating system, and other software and data for performing information processing using the main, or host, processor (CPU) within the information processing cartridge 43.

In the present implementation, the BSC 203 can be implemented as a microcontroller (e.g., a Hitachi H8 microcontroller). The BSC 203 can provide various functions, including for example: dual access (for the information processing cartridges and the CSSPs 71) to PROM 201 and EEPROM 213 for boot information
5 and a FRU-ID for the information processing cartridge; channelling communication between an information processing cartridge 43 and the service processor part of the CSSPs 71; control of power on reset (POR), system reset and externally initiated reset (XIR) to the microprocessor 192; control of the power, service-required and ready-to-remove LEDs 69; upgrading of field-upgradable firmware, via the serial interface; a
10 watchdog function for the operating system; monitoring the speed of a CPU fan 214; and communications with an EEPROM 215 and the operating system via the Xbus 200.

In the present example, the BSC 203 can be powered by a 5V service bus (SB)
15 rail as soon as a CSSP 71 and a PSU 81 are fully inserted into the midplane 171, it then turns on other DC/DC converters to provide power to the remainder of the information processing cartridge 43. A BSC reset signal can be derived from a simple conventional power on reset (POR) generator that monitors a 5V supply rail.

20 In the present example a 1MByte Flash PROM 201 can be provided for storing boot variables for OpenBoot™ PROM (OBP) and Power-On-Self-Test (POST). Further OBP variables can be stored in a second 16kByte (16kB) I2C PROM 215, accessible via the SouthBridge SM Bus port over the IC Bus 197. The PROM 215 can contain 8kByte for OBP variables and 8kByte of unused space. A 16kByte I2C
25 EEPROM 213 that is accessible via the BSC 203 can contain BSC variables and FRU-ID variables. The EEPROM is nominally divided into 8kByte for FRU-ID and 8kByte for the BSC variables. Write protection for the FRU-ID is implemented by BSC firmware. Such write protection may be carried out by, for example,

acknowledging instructions to write to the protected area, but not to carry out those write instructions.

An environmental monitor sensor 215 can be provided to monitor the CPU
5 and ambient temperatures. This sensor can be accessible via the onboard I2C bus from the BSC 203.

The information processing cartridge 43 can be powered from two, diode
commoned, 9V power supply rails 216 and 217. DC/DC converters 218 can be used
10 to provide the voltage levels required by the information processing cartridge 43. The
DC/DC converters 218 are supplied by dual 9V inputs 216, 217, individually fused
219, 220 and then diode commoned 221, 222. A 5V DC/DC converter can be turned
on as soon as the FRU is fully inserted, with the BSC 203 and required portions of the
SouthBridge 199 being powered (the 5VSB rail). A field effect transistor (FET) can
15 be used to gate off the main 5V supply to the rest of the information processing
cartridge 43. The DC/DC converter outputs and the main 5V FET can be arranged not
to turn on until the BSC 203 turns them on via a signal from the SouthBridge 199.
The SouthBridge 199 can be used so that if the BSC 203 is reset (by a watchdog
timeout or after a firmware download) the state of the DC/DC converters 218 is not
20 affected. When the remaining outputs from the DC/DC converters 218 are within
specification, a PWR_GOOD signal can be asserted low to the BSC 203.

A SouthBridge resume circuit can be operable to run from 3V3, and a simple
Zener diode dropper circuit can be used to generate 3V3 from the 5VSB supply.

25

When the FRU is inserted the inrush current can be limited, for example to
<1A, and the rate of rise can be configured not to exceed a predetermined value (e.g.,
20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to
prevent damage to the connectors and to avoid generating noise. A soft start

controller 223, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND – not shown) through the midplane 171.

5

In the present example, a processor impingement fan (processor fan) 214 is configured to run at full speed to cool the information processing cartridge 43 and the fan. The speed of the processor fan and sink can be monitored by the BSC 203, using a tachometer sense pin on the microcontroller. In the event of the fan speed falling
10 below a predetermined speed, or percentage of its nominal speed (e.g., 80%), the BSC 203 can be arranged to issue an alert. The nominal speed of the fan can be recorded as part of the BSC EEPROM contents.

The midplane connector 120 for the information processing cartridge 43 is
15 used to establish the connection between the information processing cartridge 43 and the midplane. In the present example it supports up to 84 connections (pins) that will deliver SERDES outputs 224, 225, I2C signals 226, 227, and power 216, 217. Signal connections may be made through a right-angled connector. Power connections may be made through the information processing cartridge right-angled connector. The
20 connector can be configured to facilitate hotswapping of the information processing cartridge, for example by having a low insertion force and/or guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

Interrupts to the processor 192 can be encoded using an encoded interrupt
25 vector mechanism. An I-Chip Emulator (ICE) 228 functions as an interrupt concentrator, receiving all system interrupts and encoding them as an interrupt vector according to an interrupt vector code utilisable by the processor 192. In the present example, where an UltraSPARCTM processor is used, the interrupt vector encoding may be based on a 6-bit interrupt vector code.

With reference to Figure 12, there now follows a description of an example of a combined switch and service processor (CSSP) 71. In the present example, each CSSP 71 provides the functionality of a Switch 73 and of a Shelf Service Provider, or
5 Shelf Service Processor (SSP) 74.

Figure 12 provides an overview of the functional components of the CSSP 71 including the functional components of the Switch 73 and the functional components of the SSP 74. In the present example, the components relating to the Switch 73 are
10 mounted on a Switch PCB 231, and the components relating to the SSP 75 are provided on a SSP PCB 232. It will be appreciated that such component arrangements are not compulsory for successful operation and that any other component arrangement over any number of component boards can be easily achieved using conventional component arrangement techniques.

15

Firstly, with reference to Figure 12, there follows a description of functional elements of the Switch portions 73 of a CSSP 71 as contained within the CSSP enclosure 121.

20 The midplane connector 141 on the CSSP 71 establishes the connection between the CSSP 71 and the midplane 171. In the present example, it supports up to 84 connections (pins) that will deliver SERDES outputs 265-268, I2C signals 310, 320, 321 and 322, and power 278, 279. Signal connections may be made through a 40-pair right-angled connector. Power connections may be made through a right-
25 angled connector. The connector can be configured to facilitate hotswapping of the board, for example with a low insertion force. The connector also uses guide pins to increase the ease of serviceability and prevent module misalignment during insertion.

A switch microprocessor 240 is provided, in the present example the microprocessor used is a Power PC (MPC8245) packaged in a 352pin Tape Ball Grid Array (TBGA) package. This microprocessor 240 supports between 1MB and 2GB of address space in the present example. It further includes an Embedded Programmable
5 Interrupt Controller (EPIC) that provides 5 hardware interrupts (IRQs) or 16 serial interrupts. There are 4 programmable timers with cascade mode function. DRAM memory for the processor can provided in the present example by a commodity DIMM 242. The processor 240 can be connected to a 32bit PCI bus 241, which operates at, for example, 33MHz/66MHz.

10

A clock input to the processor 240 can be provided by a clock generator (CLK) 243. The CLK 243 can include a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper
15 settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

20 In the present embodiment two switch ASICs (application specific integrated circuits) 244, 245 are provided (in the present example, BCM5632 Gigabit switch ASICs). Each ASIC can provide twelve GMII Interfaces (1 Gigabit Ethernet) (for uplinks and downlinks) and one 10Gb XGMII interface for chip-to-chip communication (bridging) 246 between the ASICs 244 and 245. Sixteen GMII 1Gb
25 'downlinks', in the form of serialized Gb Ethernet data, are provided through four quad SERDES 248-251 to allow each information processing cartridge 43 to communicate with the switch 73. Eight GMII 1Gb 'uplinks' are provided for external communication through two quad PHYs 253 and 254 (in the present example

BCM5404 ASICs) and RJ45 connectors on the rear panel 122. The ASICs 244 and 245 are configured via a PCI interface (32bit/33MHz) to the PCI bus 241.

5 A Flash PROM 256 can store a real time operating system, and management and configuration data for the microprocessor. The Flash PROM 256 in the present example can be operable to hold 8MB - 16MB of data, depending on the software required. The flash PROM 256 can be operated via an on-chip XBus 258.

Also connected to communicate with the processor 240 via the XBus 258, a
10 Real Time Clock (RTC) 259 can be provided for real-time functions with a back-up battery.

Also connected to the XBus 258 can be a UART 260 which in turn connects to a serial bus 261 for providing an asynchronous console connection from the switch 73
15 to the SSP 74 which can be accessed by the SSP.

An integrated MAC/PHY (Media Access Control/Physical) switch 271 can provides its own interface to the PCI bus 241. This MAC/PHY switch 271 can connects to a 10/100 Ethernet hub 272. The hub 272 can be operable to provide a
20 management interface to the SSP 74 and a connection from an external management network to the switch 73 and SSP 74 of a given CSSP 71. The connection from the integrated MAC/PHY device 271 to the SSP 74 can be coupled capacitively. A loopback mode can be provided by the MAC/PHY device 271 for system diagnostics.

The hub 272 can connect to an RJ45 connector 273 on the rear panel 122 of the CSSP
25 enclosure 121.

An 8kByte I2C EEPROM 262 can be used to store the FRU-ID and is accessible by the SSP portion 74 of each CSSP 71 via a serial bus 263 and the

midplane 171. The upper 2kByte of the EEPROM 262 can be configured to be write protected.

5 An I2C Redundant Control Register (RCR) 275 can be used to provide an alternate, redundant path for powering-down the CSSP 71 and Shelf Level Indicators 69 mounted on the front 57 and rear 59 panels of the shelf 41. The I2C RCR 275 can be accessible by both the SSP 73 of the CSSP 71 containing the RCR and the SSP 73 of a further CSP 71 connected via the midplane 171 via an I2C bus 276. In the present example, a device suitable for use as the RCR 275 is a Phillips PCF8574 IC.

10

With continued reference to Figure 12, there now follows a description of functional elements of the Shelf Service Provider (SSP) portion 74 of a CSSP 71 as contained within the CSSP enclosure 121 and provided on an SSP PCB 232.

15 In the present example, communication between the Switch PCB 231 and the SSP PCB 232 is facilitated by an interboard connector pair 298 and 299. It supports connections (pins) for I2C signals, 10/100 MAC/PHY output, and power. As described above, the switch PCB 231 carries the components associated with the switch, and it also carries the power, FRU-ID and environmental monitoring

20 components along with the connectors for connections to the midplane 171 and external connectors. Thus, in the present example, all SSP components requiring a connection to the midplane 171 or an external connection have signal paths routed through the connector pair 298, 299 and via the switch PCB 231 to the relevant midplane or external connectors.

25

In the present example, the SSP 74 includes a microprocessor 301 (e.g., a Power PC (MPC8245) processor) mounted on the SSP printed circuit board (PCB) 232. The processor 301 can be connected to a PCI bus 302, the present instance a 32 bit bus that operates, for example, at 33MHz/66MHz.

A clock input to the processor 301 can be provided by a clock generator (CLK) 303. The CLK 303 can comprise a configurable clock generator (not shown) implemented as a programmable clock synthesiser employing a crystal used to produce CPU clock signals. The clock frequency can be determined by jumper settings (not shown). A vectored interrupt controller (I-Chip) (not shown) and a configurable core voltage regulator module (VRM) (not shown) can be provided that operate substantially as described above with reference to the like components of Figure 11.

10

The processor 301 can be provided with a DRAM memory 305. The memory capacity can be chosen to suit the processor addressable memory space. In the present example, 8 MB of DRAM memory is provided.

15

An integrated MAC/PHY switch 306 can provide its own interface to the PCI bus 302. The MAC/PHY switch 271 can be connected to 10/100 Ethernet hub 272 via the interboard connectors 298, 299. A loopback mode can be provided by the MAC/PHY switch 306 for system diagnostics.

20

Octal UARTs (Universal Asynchronous Receiver Transmitters) 308 and 309 can be connected between the PCI bus 302 and the interboard connector pair 298, 299. The signal path can be continued from the interboard connector pair 98, 299 to serial connections 310 on the midplane connector 141 on switch PCB 231. The Octal UARTS 308, 309 can facilitate serial (I2C) communications between the SSP 74 and each of the processing cartridges 43.

25

Also connected to the PCI Bus 302 can be a dual UART (DUART) 312 that in turn can connect via the interboard connectors 298, 299 to serial bus 261 for providing an asynchronous console connection from the SSP 74 to the switch 73. The DUART

312 can also have an I2C connection to an external connector on the rear face 122 of the CSSP enclosure 121. The external connector can provide a common operating system/boot console and command port 311.

5 Connected to the processor 301 via an XBus 314 can be a Flash PROM 315. The Flash PROM 315 can store a real time operating system, and management and configuration data for the microprocessor 301. The Flash PROM 315 can be operable in the present example to hold up to 2MB of data, depending on the software required.

10 Also connected to the processor 301 via the XBus 214 can be a real time clock (RTC) 316 for real-time functions with a backup battery. The RTC 316 can also provide 8kByte of non-volatile random access memory (NVRAM), in the present instance implemented as an EEPROM. This can be used to contain information such as the FR-ID, a serial number and other FRU information.

15

 To facilitate I2C communications between the SSP 74 and the other CSSP 71, the midplane 171 and the PSUs 81, a multiplexer 318 can be provided. The multiplexer 318 can have a single I2C connection to the processor 301 and connections, via the interboard connector pair 298, 299 and the midplane connector
20 141 to both PSUs 81, the midplane 171 and the other CSSP 71.

 The processor 301 can also comprise an embedded DUART to provide a redundant serial link to the SSP 74 of the other CSSP 71. Although it would be possible to implement this link using an external DUART, the advantage of using an
25 embedded DUART is that the connection to the other CSSP is reliable and therefore likely to be functional. Where the embedded DUART link does not use the I2C Multiplexer for communications to the other CSSP, a common mode of failure for both the SSP – SSP I2C links can be avoided, it being assumed that the processor 301 is likely to be functional even if both embedded DUART channels are non-functional.

The CSSP 71 can be powered from two, diode commoned, 9V power supply rails 278 and 279. DC/DC converters 281 can be used to provide the voltage levels required by the CSSP 71. The DC/DC converters 281 can be supplied by dual 9V inputs 278, 279, individually fused 285, 286 and then diode commoned 287, 288. A soft start controller 283 can be provided to facilitate hot-insertion. A 5V DC/DC converter (I2C power regulator) 282 can be turned on as soon as the CSSP 71 is fully inserted. A 3.3V DC/DC converter can be turned when instructed, for example through SSP service software, by asserting low an appropriate signal (ON_L – not shown). The 3.3V converter can be arranged to turn on a converted for 2.5V, 1.2V, and a processor core voltage rail (Vcore) when the voltages are within an appropriate range.

When the CSSP 71 is inserted the inrush current can be limited, for example to <1A, and the rate of rise can be configured not to exceed a predetermined value (e.g., 20A/s) to provide a so-called soft start to facilitate hot-insertion. The intent is to prevent damage to the connectors and to avoid generating noise. A soft start controller 283, which controls a ramping-up of voltage levels, can be enabled when the predetermined signal (Inserted_L signal) is asserted low, this signal is on a short pin in the connector and is connected to ground (GND – not shown) through the midplane 171 until one of the supplies is removed. These circuits can be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit. A sense circuit can detect if the voltage has dropped below a threshold, for example 2.0V, as a result of a blown fuse, a power rail going down, etc. The DC/DC converters 281 can be protected against short circuit of their outputs so that no damage occurs.

The I2C regulator 282 can be powered as soon as the CPPS 71 is fully inserted into the midplane 171. This can be facilitated through short pins connected to the soft

start controller 283, which controls a ramping-up of voltage levels. The other DC/DC regulators can be turned on, for example by SSP software.

A pair of fans 290, 291 can provide cooling to the CSSP 71. The fans 290,
5 291 can be configured to run at full speed to prevent overtemperature conditions by minimizing the temperature of the internal components and the fan. The speed of the fans 290, 291 can be monitored by the SSP 74 through an environmental monitor 295 on the switch board 231. The environmental monitor 295 can be alerted in the event of the fan speed falling below a predetermined value (e.g., 80% of its nominal speed).
10 The fan can provide tachometer outputs to facilitate the measurement of fan speed.

LED indicators 137 can be provided, for example with a green power LED, an amber LED for indicating that service is required and a blue LED for indicating that the switch is ready to be removed. LED indicators integrated on 2x4 stacked RJ45
15 connectors on the rear face of the CSSP 71 can be arranged, for example, to show green continually when the link is present and flash green when the link is active.

The environmental monitor ENV MON 295 can be provided to maintain operational integrity of the CSSP 71. The ENV MON 295 can include limit values in
20 limit registers and can monitor, for example, temperature within the CSSP enclosure 72, the CSSP power rails, including the 12V, 3V3, Switch Processor Core Voltage, CSSP Processor Core Voltage and the two 9V power feed rails 278, 279 from the midplane 171. The outputs of the DC/DC converters 281 can be fed in to A/D inputs of the ENV MON 295 for Watchdog comparisons to be made to the voltage limits set
25 in the limit registers. As noted above, the ENV MON 295 can also monitor the operating speeds of the fans 290 and 291. The ENV MON 295 can communicate with the SSP 74 of both CSSPs via an I2C bus 296.

For IO to the midplane 171 shown in Figures 8A-10, the midplane connector 141 can include sixteen 1Gb Ethernet connections 265-268 from four quad SERDES 248-251 and the I2C bus lines 596.

- 5 The SSP 74 can access the I2C devices (FRU-ID EEPROM, 8-bit I/O expansion chip, and the system hardware monitor) through the midplane 171.

- For external IO, rear panel Gb Ethernet connections can be provided from the two quad PHYs 253, 254 to 2x4 stacked RJ45 connectors 139 (to give 8 uplinks).
- 10 Each port can be an independent 10/100/1000 BASE-T (auto negotiating) port. The PHY devices 253, 254 can operate in GMII mode to receive signals from the 8-Gigabit interfaces on the ASICs 244, 245.

- The Power Supply Units (PSUs) 81 can configured such that when two or
- 15 more PSUs 81 are connected in parallel in the shelf 41, failure of any one of the paralleled units shall not affect system operation. Moreover, one of the PSUs can be installed or removed from a "live" system with or without input power applied. The outputs can have overcurrent protection.

- 20 The PSU can have an I2C interface to provide power supply status via the midplane 171. The PSU can have an internal temperature sensor that reports via the I2C interface. The PSU fan speed can also be monitored and errors are reported via the I2C interface. Overvoltage and overcurrent sensors can also report via the I2C interface.

25

There now follows a description of aspects of an example of a power supply 81 with particular reference to Figure 13.

When a power supply (e.g. mains power, or UPS type protected power) is connected to the cable connector 83, transformer, regulator and rectifier circuitry 400 can operate to generate a DC output (in the present example, 12V DC) from the input (in the present example 230/240V 50Hz AC or 110V 60Hz AC).

5

In order to ensure good cooling reliability within the shelf 41, each PSU 81 can have a pair of cooling fans 402, 403 located at the rear of the PSU enclosure as described above with reference to Figure 7. As a failsafe measure against a failure of one PSU 81, the fans of each PSU 81 can be powered by both PSUs 81. Thus, in the event that one PSU 81 is non-operative to produce the DC supply for the shelf 41, as long as the other PSU 81 is operative, not only will all components of the shelf 41 be powered as normal, the fans of both PSUs 81 can continue to run.

As shown in Figure 13, this dual powering of cooling fans 402, 403 can be effected by providing a power supply line 404 from the transformer, regulator and rectifier circuitry 400 to power both fans 402, 403. Also, first and second separate power lines 410, 412 from the other PSU 81 can provide duplicate power supply to the first and second fans 402, 403 respectively. The fan 402 can thus be powered by a diode commoned supply from line 404 and a diode commoned supply from line 410. Diode protection can be provided by diodes 405 and 411 respectively. The speed of the fan 402 can be controlled by a speed controller 408. Similarly the fan 403 can be powered by a diode commoned supply from line 404 and a diode commoned supply from line 412. Diode protection can be provided by diodes 406 and 414 respectively. The speed of the fan 403 can be controlled by a speed controller 409.

25

The two speed controllers 408, 409 can in turn be controlled by a data input from each CSSP 71 received via an I2C bus connection (not shown in Figure 13). Power supply lines carrying DC power for the other FRUs of the shelf 41 are shown in Figure 13 as power line 416. All power connections to and from the PSU 81 can

connect to the midplane 171 when the PSU is inserted in the shelf 41 via the midplane connector 163. In the present example the PSU 81 connects to the shelf through a 5P/24S/6P configuration SSI-MPS compliant right angle connector 163 at the front face 147 of the PSU 81. Connectors for the I2C interface can also be provided.

5

The input power line 412 for the second fan 403 can be provided with a softstart module 413, to allow for hot insertion of the PSU 81 into the shelf 41. The softstart module 413 can be controlled, for example, by pulling a signal to ground (e.g., a "mated" input line 415). Softstart protection for the first fan 402 can be
10 provided by the main PSU softstart providing the power.

Where the two input power lines 410 and 412 are separate lines having separate softstart provision, there is no common failure mode for the backup method of powering the fans 402, 403. Thus even if the PSU 81 ceases to be operable to
15 generate the DC supply, and a component (power line or softstart module for example) fails in the supply path from the second PSU 81 to the fans 402, 403, at least one of the fans 402, 403 can remain operational as the shelf 41 still receives the cooling effect of three PSU fans.

20 In the present example, the power supply has four rear panel LED indicators 137. A blue "Ready to Remove" LED can be driven by the I2C interface and indicate that the power supply may be removed from the system. An amber "Service Required" LED can be driven by the I2C interface and indicate that the power supply is in a fault condition: any output out of range, over-temperature or shutdown. A
25 green "DC Output-OK" indicator can be driven by internal power supply circuits and show that the main 12 volt supply is functioning. The LEDs can remain lighted when individual outputs are in the current limited mode of operation. A green "AC Input-OK" indicator can be driven by internal power supply circuits and show that AC input power is within normal operating range.

With reference to Figure 14, there will now be described an example of data connectivity between the FRUs and midplane of the shelf 41. Power transmission paths are not illustrated in Figure 14. However, it will be appreciated that to facilitate
5 the maximum component redundancy of the shelf 41, each PSU 81 can independently provide power to each FRU.

In the present example each of the processing cartridges (blades) 43 connects to the midplane 171 via a pair of information signal connections (e.g. Gb Ethernet
10 links) 224, 225 and a pair of serial management signal connections 226, 227. Connections within the midplane 171 can ensure that each Ethernet link 224 is directed to a connection 265-268 from the midplane 171 to a first switch 73, and that each Ethernet link 225 is directed to a connection 265-268 from the midplane 171 to a second switch 73. Thus one Ethernet link can be established between each processing
15 cartridge 43 and the switch 73 of each CSSP 71. Further connections within the midplane 171 can ensure that each serial connection 226 is directed to a connection 310 from the midplane 171 to the first SSP 74 and that each serial connection 227 is directed to the second SSP 74. Thus one serial link can be established between each processing cartridge 43 and the SSP 74 of each CSSP 71. As mentioned earlier,
20 information signal connections other than Gb Ethernet connections (e.g. Infinband connections) could be employed in other examples.

A plurality of serial connections can connect each SSP 74 to the other. Serial lines 320, 321 can connect each SSP 74 to the midplane 171 and connections within
25 the midplane can connect the two sets of lines together. To provide a control interface from the SSPs 74 to the PSUs 81, serial lines 322 can connect each SSP 74 to the midplane 171 and connections within the midplane 171 can connect to serial lines 324 from the midplane 171 to each PSU 81.

An example of data and control connectivity of the shelf 41 to and from computer systems external to the shelf 41 when the shelf 41 is arranged for use within a multiprocessor server system such as that described above with reference to Figure 1 will be described with reference to Figure 15.

5

As summarised above with reference to Figure 14, in the present example each processing cartridge, or blade, 43 is connected to the switch 73 of each CSSP 71 by an information signal connection (e.g. a 1Gb Ethernet link) formed by a combination of links 224, 225 from the processing cartridge 43 to the midplane 171, connections
10 within the midplane 171 and links 265-268 from the midplane 171 to each switch 73.

Further, in this example a set of serial management signal connections comprising links 320, 321 and connections within the midplane 171 connect the SSP 74 of each CSSP 71 to the SSP 74 of the other CSSP 71.

15

To provide external data connectivity between the shelf 41 and an external core data network 330, in association with which all information processing performed by the processing cartridges 43 of the shelf 41 is undertaken, connections 331 can be formed between the core data network 330 and the eight 1Gb Ethernet
20 ports 139 provided on the rear panel 122 of the CSSP enclosure 121.

In the present example, the connections by means of which control and configuration of the shelf 41 are performed are entirely separate to the connections to the core data network 330. Therefore, a first external switch 335 can connect to a
25 management (I2C) port 73 of the first CSSP 71 and a second external switch 336 can connect to a management (I2C) port 273 of the second CSSP 72. As described above with reference to Figure 12, the management port 273 can provide a management network interface to both the switch 73 and SSP 74 of each CSSP 71. The external switches 335, 336 can each be connected to each of a pair of System Management

Systems (SMSs) 338, 339. The SMS is not essential to the operation of the shelf 41, but use thereof aids optimal operation of the shelf 41. In a typical multiprocessor server system a plurality of shelves 41 may be connected together via the core data network 330 under the control of a single management network utilising one set of
5 SMSs 338, 339. A set of SMSs 338, 339 may comprise a single SMS (as well as a plurality thereof). However use of at least two SMSs enables redundancy of components, therefore increasing overall system reliability.

A serial interface control 343 operable under telnet protocol control is also
10 connected to the shelf 41 in the present example. This can provide a common operating system/boot console connection to the SSP 74 of both CSSPs 71 via the RJ45 connector 311 on the rear panel 122 of each CSSP enclosure 121.

It will be appreciated from the above that a flexible and scalable modular
15 computer architecture has been described. In the described example up to 16 information processing cartridges, or blades 43, can be configured as sealed FRUs on a single shelf 41, the number of blades being chosen according to customer requirements. Each blade has its own processor and random access memory. If, for example, there is a maximum of 2Gbytes of memory per information processing
20 cartridge, and one processor per blade, 16 processors (16P) with 5.33 processors per unit height (1U) and a total of 32GB of memory per shelf can be provided.

In the present example, the shelf 41 incorporates redundant combined switch and shelf service processor modules (CSSPs) 71 and redundant power supply units
25 (PSUs) 81 separate from the blades 43. As the power supplies are carried by the shelf, the information processing cartridges can be kept compact and inexpensive. Also, as a result, they can be powered by DC power only, via the midplane 171.

Also, as mentioned earlier, the FRUs (e.g., the information processing cartridges, or blades, 43, the CSSPs 71 and the PSUs 81) can all be configured as sealed units that do not contain any internal FRUs themselves and do not contain user serviceable items. The enclosures of the FRUs can be arranged to enclose all of the
5 functional components of the FRU with only electrical connectors being externally accessible and with indicator LEDs being externally visible as well.

These factors can all contribute to keeping the cost of the FRUs low, as well as that of the overall system. The modular approach with the use of sealed modular field
10 replaceable units for providing system functionality and with non-field replaceable units designed with a minimum possible number of active components enhances reliability. Moreover, easy and rapid maintenance is facilitated in the event of a failure of a FRU by simple replacement of that FRU, further reducing the cost of ownership.

15

Thus, it will be appreciated from the above description that the provision of a rack mountable shelf, that includes power supplies, a shelf service processor and switches in modular units, for carrying a number of processing cartridges, wherein the number of processing cartridges can be chosen according to customer requirements,
20 provides a flexible and scalable computer configuration. The balancing of the load between the processors of the processing cartridges can be effected by software using conventional principles.

A configuration as described provides an easily scalable processor
25 architecture, whereby the processing power provided by a complete system based on the information processing cartridge/information processing cartridge carrier architecture can be scalable from moderate to very high capacity through the simple addition of further information processing cartridges.

As illustrated in Figure 16, an example of the external connections from a shelf 41 can be in the form of two active information signal connections (e.g., Ethernet connections) 350 and 351, two active power connections 353 and an active/standby pair of management connections 354. With regard to the management
5 connections, each connection comprises a serial connection and a network (e.g. Ethernet or Infiniband) connection. It is possible to connect to either the active or the standby connection, as the incoming signal will be internally routed to whichever management controller (CSSP) is the current master. It will be appreciated, therefore, that the connections to a shelf can be kept to a minimum. It will further be
10 appreciated from the configuration shown in Figure 16 that the system is scalable beyond a single shelf unit 41.

Figure 17 illustrates how a plurality of shelves can be configured within one (or more) racks to provide even higher processing power. Such a constellation of
15 shelves to provide a large grouping of servers is sometimes termed a “web farm” or “server farm” 360. As shown in Figure 17, the web farm comprises a plurality of shelves 41 that each carry a plurality of blades 43. Also provided are a plurality of Network Attached Storage devices (NAS) 373 for providing storage for critical data, e.g. email data storage, for the web farm. The NASs 373 are not required if there is
20 no critical data to be stored, e.g. if the web farm is operating solely to provide web caching services.

Management control of the web farm 360 can be provided through a pair of System Management Servers (SMSs) 362. Each SMS 362 can be connected to a
25 management network via a link 366 and to a management console 365. The SMSs 362 can communicate with the individual shelves 41 via a pair of management switches 364. Each shelf 41 and NAS 373 can be connected to each management switch 364 via a connection 367. Thus dual redundant management connections can be provided to each shelf 41 and NAS 373.

Flow of data to and from the web farm 360 can be provided through a pair of data switches 369. Each data switch 369 can be connected to a consumer network via a link 370. It is to be understood that the consumer network can be a larger data network to which the web farm 360 is connected. This network can be an office or corporation intranet, a local area network (LAN), a wide area network (WAN), the Internet or any other network. Connections between the data switches and the shelves 41 can be facilitated by connections 371. It is to be noted that as each shelf has its own switching capability, there is no need for each shelf 41 to be directly connected to the data switches 369. Connections can also be provided to connect the NAS units 373 to the shelves 41. The topology used for interconnection of the data switches 369, shelves 41 and NASs 373 can be any topology providing at least one connection of any length between every possible pair of units. Complex topologies arranged to minimise the maximum connection length between any two given units in the web farm can be used.

The web farm 360 comprising a plurality of shelves 41 with or without a plurality of NASs 373 can suitably be used as any or all of the entry edge server group 9, web edge server group 15 and application servers 19 described above with reference to Figure 1.

As an alternative to providing critical data storage within a NAS 373, such storage can be provided within one or more NAS cartridges fitted into one or more of the shelves 41 in place of processing cartridges 43. Another alternative is to provide a server shelf with local storage (such as a RAID array (Redundant Array of Inexpensive Disks) in place of the NAS 373.

Thus there has now been described an example of a fully configurable computing system based on a plurality of self contained field replaceable units (FRUs)

and scalable from a single processing cartridge with power supply and switching capability to a multiply redundant multiprocessor server system with full system management capability extending over a number of co-operably connected server shelves. It will of course be readily apparent to the skilled reader that many of the
5 specific features specified in the above description are in no way limiting and a variety of alternatives may be produced using only ordinary skill and common general knowledge. Non-limiting examples of example modifications which may be made to the above described system are discussed hereafter.

10 There is no limit placed on the processing cartridges as to what software they should run. Each module within a shelf or farm may run under the same operating system, or a plurality of different operating systems may be used. Examples of possible operating systems include Sun Microsystems' Solaris® OS or another UNIX™-Type OS such as Linux™, MINIX™, or Irix™, or UNIX™ or a Microsoft
15 OS such as Windows NT™, Windows 2000™, Windows ME/98/95™, Windows XP™.

It is also not necessary that each processing cartridge within a shelf or farm be configured to run the same program software. For example, individual processing
20 cartridges may be configured to execute, for example, fileserver software, mailserver software, webhosting software, database software, firewall software, or verification software.

Although it has been described above with reference to Figures 4, 8, 12, 13
25 and 14, that functionality of a switch and of a shelf service processor should be provided within a single combined switch and service processor unit, this is not essential and separate switch and shelf service processor field replaceable units may be used.

Although it has been described above that a pair of PSUs and a pair of CSSPs may be provided so as to enable dual-redundancy, further PSUs and CSSPs may be provided so as to increase FRU redundancy further, thus providing statistically higher reliability.

5

In the power supply circuitry in each of the blades and CSSPs, two voltage sense circuits may be provided after the fuses and before the diodes, to prevent a latent fault caused by a failed fuse going undetected until one of the PSUs is removed or taken offline. Such circuits may be configured to withstand an overvoltage at their inputs whilst the input they are feeding is not powered, without any leakage to the unpowered circuit.

Although it has been described above with particular reference to Figure 11 that the processing module may be based on an UltraSPARC processor, this is not limiting and any other processor having sufficient processing capacity to undertake the tasks required of a particular processing cartridge may be used. Alternative processors include, but are not limited to, Intel x86 series and compatible processors, AMD x86 compatible processors, Alpha processors and PowerPC processors. The particular example of an x86 compatible processor is described in more detail with reference to Figure 18. In Figure 18, the parts corresponding to those of the UltraSPARC based system of Figure 11 have the same reference numerals and will not be described again here. In a system based on an x86 compatible processor, the processor 378 itself communicates with the other components, including the memory 196 and PCI bus 198 via a Northbridge 379. The Northbridge 379 also includes an interrupt controller, so no separate interrupt concentrator is required. The other components of the processing cartridge could be substantially the same as for the UltraSPARC based system described above.

Although it has been described above that each information processing cartridge comprises a single microprocessor, this is not a limiting case as each or any of the information processing cartridges may have more than one microprocessor arranged to share common storage resources to operate synchronously (in lockstep) or
5 asynchronously. Also, it is not necessary that all information processing cartridges inserted into a shelf at a given time are identical, rather a variety of different blade architectures may be used simultaneously.

The provision of the functions of both Switch and Shelf Service Processor
10 within a single FRU in the present example provides a facility within a single shelf 41 for dual redundancy in both functions in fewer different FRUs. As will be appreciated, there is no restriction that these two functions must be provided within a single FRU and division of the two functions into separate FRUs would present no difficulty to the skilled addressee.

15

The backup power provision for the cooling fans of the PSUs is described above with reference to Figure 13. Although it is described that the backup power supply to each fan should be independent of the other, if the risk of common mode failure of backup power supply is judged to be low, or is of low importance, then the
20 backup supply to each fan may be provided in common with the backup supply to all other fans.

Also, an information processing module forming a field replaceable server blade can include a processor and memory can be configured by means of software,
25 firmware or hardware to provide a special purpose function. By way of examples only, an information processing module can be configured to perform the function of one or more of a firewall, or a load balancer, encryption and/or decryption processing, an interface to a secure network, e.g. a virtual private network (VPN), a specialized switch with wide area network (WAN) connectability.

Also, a storage blade may be provided. The storage blade can be configured to be mountable in a server blade receiving location in a blade server carrier. The storage blade can comprise storage blade connectors configured for interconnecting
5 with carrier connectors on the server blade carrier, whereby the storage blade is interchangeable with a server blade. A carrier, or shelf, for a server system can be arranged with a plurality of blade receiving locations for receiving blades, wherein the blades can be storage blades or information processing blades. The server system can be self configuring on receipt of the blades according to the type of blade received in
10 each said location. To achieve the blade service controller in each blade can be operable to communicate with a shelf service provider to perform said configuring.

The use of Network Attached Storage (NAS) devices has been described above. Addressing of the NAS devices is achieved over the information signal
15 connections (i.e. over the Ethernet or Infiniband connections) provided internally and externally of the shelves. This provides a very effective implementation of bulk storage which can be accessed by applications running on the individual modules (e.g. the server blades) within the shelf. However, in examples where there is a very high usage of such storage, this can have an impact on network performance.

20

Accordingly, an embodiment of the present invention seeks to reduce the impact on the external information processing network through the use of dedicated external storage connections. In particular, Figure 19 illustrates a modified combined switch and service processor to replace the combined switch and service processor of
25 Figure 12.

The embodiment described with reference to Figure 19 employs dedicated storage connections employing, by way of example, Fiber Channel format connections as a Fiber Channel Arbitrated Loop (FC-AL). Fiber Channel is a channel/network

standard that can be implemented over fiber or copper and can operate as a generic transport mechanism.

In this embodiment, additional connections are provided over the midplane
5 between the storage blades 43 and the CSSP 71. As shown in Figure 19, sixteen individual Fiber Channel connections 410 are provided on the midplane connector 141. Each of the sixteen connections 410 (0:15) relates to a respective blade server. Each of those connections can be connected via a respective differential path 412 to a respective port bypass circuit 414. An example of a port bypass circuit is a HDMP-
10 0421 circuit manufactured by Agilent Technologies. The port bypass circuits can be connected to each other to form the Fiber Channel Arbitrated Loop (FC-AL) 416. Such circuits are typically used to provide loops that are continuously on in hard disk arrays. However, in the present example, such circuits are used to form a loop with respect to the blades. Each port bypass circuit 414 is a 2:1 multiplexer array with two
15 modes of operation: "inloop" and "bypassed". In the "inloop" mode, the loop is considered to go in and out of the server blade connected to the port bypass circuit 414 concerned. In the "bypassed" mode, the server blade 43 is either absent or non-functional and the loop bypasses the server blade 43 concerned. The port bypass circuits 414 could each be implemented in a separate integrated circuit, or
20 alternatively a plurality of them could be combined in a single, larger integrated circuit.

At either end, the loop 416 is connected to a GigaBit Interface Converter (GBIC) 418 for converting between electrical signals used in the loop 416 and optical
25 signals passed by a path 420 to an optical connection 422 provided on the I/O panel connector. An optical connector can be included on the rear face of the CSSP 71. Conventional external storage devices (not shown) can then be accessed via the optical connection 422.

As an alternative, or in addition, to the optical connection 422 on the I/O panel connector, the GBICs 418 can be omitted, or bypassed, respectively, and a 9-pin electrical connector can be provided for each end of the Fiber Channel loop 416. A nine-pin electrical connector can be provided on the rear face of the CSSP 71.

5

Figure 20 illustrates a modified server blade to replace the server blade shown in Figure 11 for use with an embodiment of the invention.

It will be noted that in addition to the circuitry shown in Figure 11, the server blade shown subsystem shown in Figure 20 also includes first and second Fiber Channel interfaces 430 and 431. These generate Fiber Channel format signals to be passed via connections 432 and 433, respectively, on the midplane connector 120 to switches A and B, respectively. The Fiber Channel interfaces 430 and 431 can be conventional Fiber Channel interfaces (e.g., as defined by the American National Standards Institute).

15

It will be appreciated that although the use of Fiber Channel is described, other storage data formats can be used.

20

In the present instance, the transmission of signals over the midplane via the midplane connectors of the server blades and the combined switch and service processor is via electrical connections. It will be appreciated that the provision of the separate Fiber Channel electrical pathways will mean that the midplane connectors of the server blades and the switch and service processor require specific connections for these signals. An appropriately sized connector or a set of connectors could be used. Also, appropriate connection pathways on the midplane need to be provided. However, an embodiment of the present invention has the effect of freeing up bandwidth for non-storage activities over both the internal and external information networks (i.e. the Ethernet or Infiniband) connections. Moreover, the use of a

25

structure as described, by way of example, with reference to Figures 19 and 20 enables the connection of direct-attached storage (DAS) and/or connection to a storage area network (SAN) over the external Fiber Channel connection.

CLAIMS

1. A computer system comprising:
at least one information processing module including a processor, memory, at
least one module information connection for communicating information
5 signals, and at least one module storage connection for passing data to and
from external storage; and
a carrier comprising at least one module receiving location configured to
receive an information processing module, the module receiving location
comprising at least one carrier information connection operable to interconnect
10 with a corresponding module information connection of a received
information processing module and at least one carrier storage connection
operable to interconnect with a corresponding module storage connection of a
received information processing module,
the carrier comprising:
15 at least one switch connectable to at least one carrier information connection
of at least one module receiving location and operable to connect the
information processing module to an external information network; and
at least one port bypass controller connectable to at least one carrier storage
connection of at least one module receiving location and operable to connect
20 the information processing module to external storage.
2. The computer system of claim 1, comprising a switch module removably
mountable in the carrier, the switch module comprising at least one switch and
at least one port bypass controller.
- 25 3. The computer system of claim 1 or claim 2, wherein the information
processing module is removably receivable within the carrier.

4. The computer system of any preceding claim, wherein the carrier is configured to receive a plurality of information processing modules in respective module receiving locations.
- 5 5. The computer system of claim 4, wherein each switch is connectable to a carrier information connection at each module receiving location and is operable selectively to connect each received information processing module to an external information network.
- 10 6. The computer system of claim 5, wherein each information processing module comprises two module information connections, each module receiving location comprises two carrier information connections operable to interconnect with a respective module information connection of a received information processing module and the carrier includes two switches, each
15 connectable to a respective carrier information connection of at least one module receiving location to provide redundant information signal connections.
7. The computer system of claim 6, wherein each switch is connectable to a
20 respective carrier information connection at each module receiving location.
8. The computer system of any of claims 4 to 7, wherein a plurality of port bypass controllers are interconnected to form at least one chain of port bypass controllers, each port bypass controller being connectable to a respective
25 carrier storage connection at each module receiving location, whereby a plurality of information processing modules received in respective receiving locations is connectable to the external storage.

9. The computer system of claim 8, wherein each information processing module comprises two module storage connections, each module receiving location comprises two carrier storage connections operable to interconnect with a respective module storage connection of a received information processing module and the carrier includes a first and a second plurality of port bypass controllers interconnected to form first and second chains of port bypass controllers, respective port bypass controllers of the first and second chains being connectable to a respective carrier storage connection at each module receiving location to provide redundant connections of information processing modules received in respective receiving locations to external storage.
10. The computer system of claim 9 when dependent upon claim 7, comprising first and second switch modules removably mountable in the carrier, each switch module comprising a respective switch and a respective chain of port bypass controllers.
11. The computer system of any preceding claim, wherein the storage connections comprise Fiber Channel connections.
12. The computer system of any preceding claim, wherein the information connections comprise Ethernet connections.
13. The computer system of any of claims 1 to 12, wherein the information connections comprise Infiniband connections.
14. The computer system of any preceding claim, wherein each information processing module comprises a service controller and the carrier comprises at least one service processor, each service processor being connectable to

service controllers of the plurality of information processing modules via respective management signal connections.

15. The computer system of claim 12, wherein the carrier comprises two service
5 processors, a first service processor being operable as a master service
processor and a second service processor being operable as a backup service
processor.
16. The computer system of claim 14 when dependent upon claim 10, wherein
10 each switch module further comprises a respective service processor to form a
combined switch and service processor module.
17. The computer system of any preceding claim, wherein the carrier comprises at
least one power supply connectable to each information processing module
15 received in a carrier receiving location via respective power connections.
18. The computer system of claim 17, wherein the carrier comprises a plurality of
power supplies, each power supply being connectable to each information
processing module received in a carrier receiving location via respective power
20 connections.
19. The computer system of claim 17 or claim 18, wherein each power supply is
removably receivable in the carrier.
- 25 20. The computer system of any preceding claim, wherein the carrier is a rack
mountable shelf.
21. The computer system of any preceding claim, wherein the information
processing unit comprises an enclosure enclosing the processor and memory.

22. The computer system of any preceding claim, wherein the processing unit further comprises an internal storage device.
- 5 23. The computer system of any preceding claim wherein the information processing module is module is a blade server.
24. A server blade comprising at least one processor, memory, at least one blade information connection for communicating information signals, and at least
10 one blade storage connection for passing data to and from external storage.
25. The server blade of claim 24 comprising a plurality of blade information connections for communicating information signals, and a plurality of blade storage connections.
- 15 26. The server blade of claim 24 or claim 25, comprising a Fiber Channel interface for each blade storage connection.
27. The server blade of any of claims 24 to 26, comprising an Ethernet interface
20 for each blade information connection.
28. The server blade of any of claims 24 to 26, comprising an Infiniband interface for each blade information connection.
- 25 29. The server blade of any of claims 24 to 28 comprising a service controller operable to provide system management functions for the blade server.
30. The server blade of claim 24 to 29, comprising an internal storage device.

31. The server blade of any of claims 24 to 30 configured to be hot-swappable.
32. The server blade of any of claims 24 to 31 comprising an enclosure enclosing the processor and memory, the blade connections being accessible externally to the enclosure.
33. A switch module for a carrier that includes at least one receiving location for receiving an information processor module that includes at least a module information connection for communicating information signals, and a module storage connection for passing data to and from storage, the switch module comprising:
- a switch connection for each receiving location operable to connect a received information processing module to an external information network; and
- a port bypass controller for each receiving location operable to connect a received information processing module to external storage.
34. The switch module of claim 33, wherein each port bypass controller is interconnected to form a chain of port bypass controllers.
35. The switch module of any of claim 33 or claim 34, comprising an enclosure.
36. The switch module of any of claims 33 to 35, removably receivable in the carrier.
37. The switch module of any of claims 33 to 36, further comprising a service processor to form a combined switch and service processor module.
38. A carrier for a computer system, the carrier comprising receiving locations for a plurality of information processing modules, at least one switch module and

an interconnection member, the interconnection member interconnecting information signal channels and storage signal channels between the information processing modules and the switch.

- 5 39. The carrier of claim 38, wherein the interconnection member is a midplane carrying connectors for each of the information processing and switch modules, and interconnecting pathways.
40. The carrier of claim 38 or claim 39, configured to receive first and second
10 switch modules.

1/24




Fig. 1

2/24




Fig. 2




Fig. 3

4/24





Fig. 4

6/24



8/24




Fig. 8A

10/24




Fig. 8C

11/24




Fig. 9A




Fig. 9B




Fig. 9C

12/24





Fig. 10

14/24



SUBSTITUTE SHEET (RULE 26)

15/24




Fig. 12

16/24




Fig. 13




Fig. 14

18/24




Fig. 15

19/24




Fig. 16

20/24




Fig. 17

21/24






Fig. 18

22/24



continued on sheet 23/24

23/24



24/24




Fig. 20

